

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (Currently Amended) An apparatus for bandwidth management, comprising:

a plurality of load shapers configured to:

maintain a local bandwidth management table comprising a local token count for each of a plurality of classes of source entities;

receive a data packet from a source entity belonging to one of the plurality of classes;

transmit the data packet over a multiplexed communication path if the local token count for the class of the source entity is at least one; and

decrement the local token count for the class of the source entity in the local bandwidth management table in response to the transmission; and

a Bandwidth Management Controller configured to:

maintain a centralized bandwidth management table comprising a base token count for each of the plurality of classes of source entities, wherein a minimum bandwidth is reserved for each of the plurality of classes of source entities and the base token count increases at a rate corresponding to the minimum bandwidth; and wherein:

the plurality of load shapers is further configured to request a token for the class of the source entity from the Bandwidth Management Controller in response to the transmission; and

the Bandwidth Management Controller is further configured to respond to the request if the base token count for the class of the source entity is at least one by:

providing a[[.]] token and decrementing the base token count for the class of the source entity.

2. (Currently Amended) The apparatus of claim 1, wherein:

the centralized bandwidth management table further comprises a standby token count for each of the plurality of classes of source entities; and

the Bandwidth Management Controller is further configured to respond to the request if the base token count for the class of the source entity is zero and the standby token count for the class of the source entity is at least one by:

providing a[[.]] token and decrementing the standby token count for the class of the source entity.

3. (Previously Presented) A method of bandwidth management, comprising:

maintaining a centralized bandwidth management table comprising a base token count for each of a plurality of classes of source entities, wherein a minimum bandwidth is reserved for each of the plurality of classes of source entities and the base token count increases at a rate corresponding to the minimum bandwidth;

maintaining a local bandwidth management table comprising a local token count for each of the plurality of classes of source entities;

receiving a data packet from a source entity belonging to one of the plurality of classes;

transmitting the data packet over a multiplexed communication path if the local token count for the class of the source entity is at least one;

decrementing the local token count for the class of the source entity in the local bandwidth management table in response to the transmission;

requesting a token for the local token count for the class of the source entity in response to the transmission; and

if the base token count for the class of the source entity is at least one, providing the requested token and decrementing the base token count for the class of the source entity.

4-16. (Canceled)

17. (Previously Presented) The apparatus of claim 2, wherein the Bandwidth Management Controller is further configured to:

linearly increase the standby token count for each of the plurality of classes of source entities when the communication path is not congested, and

exponentially decrease the standby token count for each of the plurality of source entities when the communication path is congested.

18. (Previously Presented) The apparatus of claim 1, wherein the local token count for each of the plurality of classes of source entities has a maximum count of two tokens.

19. (Previously Presented) The apparatus of claim 1, wherein the plurality of load shapers is further configured to maintain a count of outstanding requests for tokens.

20. (Previously Presented) The method of claim 3, wherein:

the centralized bandwidth management table further comprises a standby token count for each of the plurality of classes of source entities;

and further comprising:

if the base token count for the class of the source entity is zero and the standby token count for the class of the source entity is at least one, providing the requested token and decrementing the standby token count for the class of the source entity.

21. (Previously Presented) The method of claim 20, further comprising:

linearly increasing the standby token count for each of the plurality of classes of source entities when the communication path is not congested; and

exponentially decreasing the standby token count for each of the plurality of source entities when the communication path is congested.

22. (Previously Presented) The method of claim 3, wherein the local token count for each of the plurality of classes of source entities has a maximum count of two tokens.

23. (Previously Presented) The method of claim 3, further comprising maintaining a count of outstanding requests for tokens.

24. (Previously Presented) A computer program product for bandwidth management, the computer program product embodied on a tangible computer readable medium comprising:

computer code for maintaining a centralized bandwidth management table comprising a base token count for each of a plurality of classes of source entities, wherein a minimum bandwidth is reserved for each of the plurality of classes of source entities and the base token count increases at a rate corresponding to the minimum bandwidth;

computer code for maintaining a local bandwidth management table comprising a local token count for each of the plurality of classes of source entities;

computer code for receiving a data packet from a source entity belonging to one of the plurality of classes;

computer code for transmitting the data packet over a multiplexed communication path if the local token count for the class of the source entity is at least one;

computer code for decrementing the local token count for the class of the source entity in the local bandwidth management table in response to the transmission;

computer code for requesting a token for the local token count for the class of the source entity in response to the transmission; and

computer code for, if the base token count for the class of the source entity is at least one, providing the requested token and decrementing the base token count for the class of the source entity.

25. (Previously Presented) The computer program product of claim 24, wherein:

the centralized bandwidth management table further comprises a standby token count for each of the plurality of classes of source entities;

and further comprising:

computer code for, if the base token count for the class of the source entity is zero and the standby token count for the class of the source entity is at least one, providing the requested token and decrementing the standby token count for the class of the source entity.

26. (Previously Presented) The method of claim 25, further comprising:

computer code for linearly increasing the standby token count for each of the plurality of classes of source entities when the communication path is not congested; and

computer code for exponentially decreasing the standby token count for each of the plurality of source entities when the communication path is congested.

27. (Previously Presented) The method of claim 24, wherein the local token count for each of the plurality of classes of source entities has a maximum count of two tokens.

28. (Previously Presented) The method of claim 24, further comprising maintaining a count of outstanding requests for tokens.

29-32. (Canceled)

33. (New) The apparatus of claim 1, wherein the apparatus is a multi-processor system comprising:

- a plurality of processor units, wherein each of said plurality of processor units is associated with a respective one of said plurality of load shapers;
- a bus interconnecting each of said plurality of processor units;
- a display device connected to said plurality of processor units and configured to display data from each of said plurality of processor units.

34. (New) The method of claim 2, wherein said method of bandwidth management is performed in a multi-processor system comprising a plurality of processor units; and

- wherein said local bandwidth management table is one of a plurality of local bandwidth management tables each being associated with a respective one of said plurality of processor units.

35. (New) The method of claim 3, wherein the centralized bandwidth management table further comprises a standby token count for each of the plurality of classes of source entities, the method further comprising:

- responding to the request if the base token count for the class of the source entity is zero and the standby token count for the class of the source entity is at least one by:

- providing a token and decrementing the standby token count for the class of the source entity.

36. (New) The method of claim 34, further comprising:

linearly increasing the standby token count for each of the plurality of classes of source entities when the communication path is not congested, and

exponentially decreasing the standby token count for each of the plurality of source entities when the communication path is congested.